

REMARKS

Claims 1-39 are currently active.

The Examiner has rejected Claims 1-11 as being unpatentable over Fridella in view of APA and further in view of Edsall. Applicants respectfully traverse this rejection.

Referring to Fridella, there is disclosed management of file-modification time attributes in a processor fileserver system. Fridella teaches a network file server includes multiple data mover computers 115, 116 and 117, each of which manages a respective filesystem. Each data mover computer also functions as a fileserver for servicing client requests for access to the filesystems. Each data mover computer has a respective port to a data network 111 having a number of clients including workstations. The network file server includes a cached disc array 114. The clustering of the data movers as a front end to the cached disk array provides parallelism and scalability. The data movers may communicate with each other over a dedicated dual-redundant ethernet connection 118. See paragraphs 26 and 27. In the network file server, the locking information for each filesystem is managed exclusively by only one of the data movers. This exclusive relationship is referred to by saying each filesystem has a respective data mover that is the owner of the filesystem. The

owner of a filesystem is said to be primary with respect to the filesystem, and other data movers are said to be secondary with respect to the filesystem. See paragraph 28.

Each client may access any of the file systems to any one of the data mover computers but if the data mover computer servicing the client does not own the filesystem to be accessed, then a lock on at least a portion of the filesystem to be access must be obtained from the data mover computer that owns the filesystem to be accessed. It is possible for a write operation to change the attributes of a file. When a write operation will change the metadata of a file, the metadata must be managed in a consistent fashion, in order to avoid conflict between the data mover owning the file, and the data mover performing the write operation. See paragraphs 29 and 30. Fridella teaches it is desirable for some of the filesystem metadata to be cached only on the primary data server, and some of the filesystem metadata to be cached on the primary and secondary data movers. See paragraph 32.

Fridella teaches it is possible for the secondary data movers to update the file-modification time attribute concerning a file in a consistent fashion without always accessing the primary data mover clock. The clocks of the primary and secondary data mover need not be synchronized. The secondary clocks cannot simply be used to set the file-modification time attributes, because the clock skew between the multiple secondary data movers writing to the same file would violate the second consistency requirement. On the other hand, the primary

clock cannot simply be used unless the file-modification time is updated for each asynchronous write. It is possible though for a secondary data mover to update the file-modification time attribute in a consistent fashion using a hybrid method that comprises the file-modification time attribute based on the clock of the primary data mover and a timer of the secondary data mover. See paragraph 36.

When a server performs an asynchronous write for a client, the server returns an updated file as a modification time attribute. If the server is the primary server, the updated-file-modification time can simply be the time of its local clock. If the server is the secondary server, then the updated file modification time is the sum of the local timer and a local value of the primary clock having been stored in local memory of the secondary server. When the secondary server performs a second asynchronous write to the filesystem, it computes an updated file-modification time by adding the stored clock time and the present value of the timer, and returns the file-modification time to the client. See paragraphs 38 and 39.

Fridella teaches the file management protocol between the primary server and the secondary server is designed to permit the exchange of file metadata between primary and secondary servers that cache the file metadata. This protocol eliminates the need for the

secondary to communicate with the primary every time that the secondary response to a read or write request.

It is clear from the above description, there is no teaching or suggestion whatsoever of "the servers utilizing a striped filesystem for storing data," as found in Claim 1 of applicants. Furthermore, the focus of Fridella is to update the file-modification time attribute of the secondary data movers without always accessing the primary data mover clock. This is the key and the focus of the architecture taught by Fridella.

Referring to APA, there is taught a PVFS system, where data is striped among multiple servers through an additional filesystem layer built on top of a normal filesystem. The PVFS updates to the various striped files and the resulting filesystem are not coordinated very closely, and operations that deal with global file properties, such as the file length, are implemented very expensively, or via approximations that may cause application errors. PVFS filesystems return the modification time for a file via a similar procedure to that returning the file length: they check with all servers and return the largest modification time field. Since the different servers have clocks that differ by no less than small numbers of microseconds, it is possible for a write to be performed at the server responsible for one stripe that happens to have the furthest advanced clock, and then perform a write to another server with an older clock, with the result that the second write does not advance the system wide file

modification time. Having two versions of the file with the same modification time may cause incorrect behavior by protocols like NFS that use modification times as version numbers. Because of these problems, PVFS filesystems are unsuitable for export over a network with NFS. This is the context of APA.

There is no teaching or suggestion in Fridella how to modify PVFS systems to overcome these problems so they are suitable for NFS. Furthermore, applicants do not suggest they discovered NFS, nor do applicants suggest they discovered striping, but applicants do take the position that they were the first to be able to apply a striped filesystem for storing data with a cluster of NFS servers. Simply identifying Fridella that has a plurality of NFS servers that have nothing at all to do with striping, and the teachings of APA that specifically identify problems that make PVFS as unsuitable with NFS cannot arrive at applicants' claimed invention.

It is respectfully submitted that the Examiner is using hindsight to arrive at applicants' claimed invention. The Examiner is using the limitations of Claim 1 as a roadmap to find the different limitations in the various prior art references, and having found the limitations, concluding that applicants' claimed invention is arrived at. This is not patent law.

Moreover, the teachings the Examiner is relying upon cannot be taken out of the context in which they are found. The context of APA is striping as it is used in PVFS systems where data is striped among multiple servers through an additional filesystem layer built on top of a normal filesystem. In terms of the timing, different servers have different clocks that differ by no less than small numbers of microseconds and with the result that later writes do not advance a systemwide file modification time. The teachings of Fridella are in regard to an NFS server system where the secondary data movers can update the file modification time attribute in a consistent fashion without always accessing the primary data mover clock. These two contexts have nothing at all to do with each other. It is a total unknown to applicants how Fridella would somehow or other be modified to use the teachings of APA to arrive at applicants' claimed invention. Because of the context of the timing in APA is totally different than Fridella, and the context of APA requires a different architecture totally, using an additional filesystem layer built on top of a normal filesystem, which is not found in Fridella, it would require significant research and development to take the striping taught in the context of APA and apply it to the NFS system taught by Fridella. This only further supports the nonobviousness and uniqueness of applicants' claimed invention.

Patent law also requires that there must be some teaching or suggestion in the references themselves to combine the teachings the Examiner is relying upon to arrive at applicants' claimed invention. Here, there is no such teaching or suggestion. There is no

reason why one skilled in the art would look at PVFS filesystem that utilizes an additional filesystem layer built on top of a normal filesystem and apply it to a totally different NFS system. Neither has any motivation for or need of any features of the other. In fact, it would be suggested that simply because one is an NFS system and the other utilizes a PVFS filesystem, they are incompatible and one skilled in the art would not look to either one to combine them to arrive at applicants' claimed invention.

Referring to Edsall, there is disclosed an apparatus and method for a scalable network attached storage system. Edsall teaches a NAS system having a scalable architecture which includes a load balancer 12, one or more termination nodes, one or more fileserver nodes, one or more disk controller nodes, and a plurality of discs 20. A switching fabric 22 is provided to interconnect the termination nodes, the fileserver nodes, and the disk controller nodes. The system is connected to a network through a standard network interconnect. See paragraph 20. A load balancer receives requests to access files stored on the system from users on the network. The main function performed by the load balancer is to balance the number of active connections among the one or more termination nodes. The load balancer dynamically assigns user connections so that no one termination node becomes a bottleneck due to handling too many connections. See paragraph 21. The load balancer attains a current load of each of the available termination nodes in the system. The termination node with the smallest current load is that identified. The new connection is assigned to the termination

node with the smallest load. The termination nodes each perform a number of functions. The termination nodes terminate connection requests received to the load balancer from clients over the network. Termination involves the conversion or translation of the upper layer protocols, such as NFS, into the communication protocol used by the switching fabric.

As is clear from the above description, the architecture taught by Edsall is distinct and different from the architectures taught by Fridella and APA. Edsall does not teach or suggest the limitation of the servers utilizing a stripe file system for storing data as found in Claim 1. It is also respectfully pointed out that while the Examiner has cited Edsall as part of the applied art of record in regard to the rejection of Claims 1-11, nowhere can applicants find any discussion in the Office Action of why the Examiner cited Edsall in regard to Claims 1-11.

Accordingly, the applied art of record does not teach or suggest the limitations of Claim 1. Claims 2-11 are dependent to parent Claim 1 and are patentable for the reasons Claim 1 is patentable.

Claim 14 is patentable for the reasons Claim 1 is patentable. Claims 15-19 and 29-35 are dependent to parent Claim 14 and are patentable for the reasons Claim 14 is patentable.

The Examiner has rejected Claims 12 and 13 as being unpatentable over Fridella in view of APA and further in view of Edsall. Applicants respectfully traverse this rejection. It is here that the Examiner's comments as to why Edsall has been cited are stated. The Examiner cites Edsall because it teaches to use a network utilizing gigabit ethernet switches to connect servers, nodes and disk controller nodes. Claims 12 and 13 are dependent to parent Claim 1 and are patentable for the reasons Claim 1 is patentable. As explained above, Edsall does not teach or suggest the limitation of the servers utilizing a served filesystem for storing data. Moreover, Edsall does not explain how one would use a network of switches to modify the teachings of Fridella, which have nothing at all to do with striping, to somehow or other accommodate striping. There must be some teaching or suggestion in the applied art of record to combine the teachings the Examiner is relying upon to arrive at applicants' claimed invention, and here, there is none. There is no reason why one skilled in the art would take switches taught by Edsall having nothing at all to do with striping, and somehow or other modify the architecture of Fridella, which does not teach or suggest the need at all for switches to connect servers, nodes and disk controller nodes in such a way as to accommodate striping. Accordingly, Claims 12 and 13 are patentable over the applied art of record.

The Examiner has rejected Claims 20, 21, 23 and 25 as being unpatentable over Fridella in view of APA and further in view of Cheng. Applicants respectfully traverse this

rejection. The Examiner cites Cheng solely for the teaching of queuing the requests. A review of Cheng shows that it has nothing at all to do with striping let alone striping for data storage. Cheng does not even teach a cluster of NFS servers. It is respectfully submitted that once again, the Examiner is simply looking for a limitation that is found in applicants' claims, in the prior art, and having found it, concludes that it would be obvious to use with the other references the Examiner cites to arrive at applicants' claimed invention. This completely ignores that there is no teaching or suggestion in the applied art of record, and certainly not in Fridella or APA for the need to queue requests. Accordingly, the applied art of record does not teach or suggest the limitations of Claim 14. Claims 20, 21, 23 and 25 are dependent to parent Claim 14 and are patentable for the reasons Claim 14 is patentable.

Applicants note that there is no formal rejection against Claim 36, only a brief comment by the Examiner that the rejection against Claim 36 is based on the same rationale as the rejection of Claim 14. While Fridella does teach meta-data, applicants do not suggest that they discovered meta-data. Claim 36 has the limitations of "receiving a file create request at a meta-data server from the network element; allocating an inode number for the file at the meta-data server; making create calls to input output secondaries to mark the file as allocated by the input output secondaries". The applied art of record does not teach or suggest these limitations.

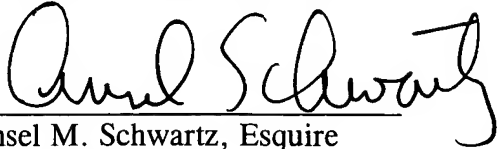
The Examiner has rejected Claim 37 as being unpatentable over Fridella in view of APA and further in view of Schmuck. Applicants respectfully traverse this rejection. Schmuck teaches a parallel filesystem with a method for using tokens for locking modes. Claim 37 has the limitations of "removing a filename of the file from a parent directory by the meta-data server at the meta-data server; putting the file on a file delete list by the is one data server at the meta-data server". It is submitted that neither Fridella nor APA, nor Schmuck teach or suggest these limitations.

The Examiner has rejected Claims 38 and 39 as being unpatentable over Fridella in view of APA and Schmuck and Ross. Applicants respectfully traverse this rejection. Claim 38 has the limitation of reading the stripes of the file from each disk element having the stripes. Claim 39 has the limitation of a plurality of servers having a stripe of the VFS, and the limitation of receiving NFS write requests for a file at a network element. Schmuck does not teach or suggest anything at all about striping, and thus does not add anything to the teachings of Fridella and APA in regard to the claimed invention. Ross discloses an implementation and performance of a parallel filesystem for high-performance distributed applications. It has to do with PVFS, like that of APA. It does not add anything at all to the teachings of APA and how to somehow or other use striping in regard to an NFS system. Accordingly, Claims 38 and 39 are patentable over the applied art of record.

In view of the foregoing remarks, it is respectfully requested that the outstanding rejections and objections to this application be reconsidered and withdrawn, and Claims 1-39, now in this application be allowed.

Respectfully submitted,

MICHAEL L. KAZAR, ET AL.

By 

Ansel M. Schwartz, Esquire

Reg. No. 30,587

One Sterling Plaza

201 N. Craig Street

Suite 304

Pittsburgh, PA 15213

(412) 621-9222

Attorney for Applicants

